

CLAIMS

What is claimed is:

1. A fault-tolerant server comprising:
 - (a) a communications link;
 - (b) a first computing element in electrical communication with the communications link, the first computing element providing a first instruction to the communications link;
 - (c) a second computing element in electrical communication with the communications link, the second computing element providing a second instruction to the communications link;
 - (d) a first local input-output (I/O) subsystem in electrical communication with the first computing element and the communications link; and
 - (e) a second local I/O subsystem in electrical communication with the second computing element and the communications link,

wherein at least one of the first local I/O subsystem and the second local I/O subsystem compares the first instruction and the second instruction and indicates a fault of at least one of the first computing element and the second computing element upon the detection of a miscompare of the first instruction and the second instruction.

2. The fault-tolerant server of claim 1 wherein each computing element further comprises a respective Central Processing Unit (CPU) and a respective local mass storage device.
3. The fault-tolerant server of claim 2 wherein the communications link further comprises a respective switching fabric in electrical communication with the CPU and at least one of the first local I/O subsystem and the second local I/O subsystem.
4. The fault-tolerant server of claim 1 further comprising a priority module to assign a priority to each respective computing element.

5. The fault-tolerant server of claim 4 wherein each local I/O subsystem further comprises I/O fault-tolerant logic to determine whether at least one of the first computing element and the second computing element is faulty based on the priority.
6. The fault-tolerant server of claim 1 wherein each local I/O subsystem further comprises I/O fault-tolerant logic to determine whether the first I/O instruction and the second I/O instruction are substantially equivalent.
7. The fault-tolerant server of claim 6 wherein each I/O fault-tolerant logic comprises a comparator.
8. The fault-tolerant server of claim 6 wherein each I/O fault-tolerant logic further comprises a buffer to hold at least one of the first I/O instruction and the second I/O instruction from at least one of the CPUs.
9. The fault-tolerant server of claim 1 further comprising a voter delay buffer to store at least one of the first instruction and the second instruction upon a miscompare of the first instruction and the second instruction.
10. The fault-tolerant server of claim 1 further comprising a delay module in electrical communication with the local I/O subsystem to delay transmission of at least one instruction to the local I/O subsystem.
11. The fault-tolerant server of claim 1 wherein the first computing element and the second computing element further comprise a 1U rack-mount motherboard.
12. The fault-tolerant server of claim 1 wherein each respective local I/O subsystem is located on a same motherboard as the respective computing element.

13. A method for a first computing element and a second computing element to execute in lockstep in a fault-tolerant server, the method comprising the steps of:
- (a) establishing communication between the first computing element and a communications link;
 - (b) establishing communication between the second computing element and the communications link;
 - (c) transmitting, by the first computing element, a first instruction to the communications link;
 - (d) transmitting, by the second computing element, a second instruction to the communications link; and
 - (e) comparing, by at least one of a local input-output (I/O) subsystem of the first computing element and a local I/O subsystem of the second computing element, the first instruction and the second instruction and indicating a fault of at least one of the first computing element and the second computing element in response thereto.
14. The method of claim 13 further comprising the step of transmitting a stop command to each computing element when the first instruction does not equal the second instruction.
15. The method of claim 13 further comprising detecting an error introduced by the communications link.
16. The method of claim 13 further comprising assigning a priority to each respective computing element.
17. The method of claim 16 further comprising determining whether at least one of the first computing element and the second computing element is faulty based on the priority.

18. The method of claim 16 further comprising determining whether the first I/O instruction and the second I/O instruction are substantially equivalent.
19. The method of claim 13 further comprising storing at least one of the first I/O instruction and the second I/O instruction from at least one of the computing elements for a predetermined amount of time.
20. The method of claim 13 further comprising storing at least one of the first instruction and the second instruction upon a miscompare of the first instruction and the second instruction.
21. The method of claim 13 wherein the transmitting of the first instruction and the transmitting of the second instruction to the communications link occur simultaneously.
22. An apparatus for enabling a first computing element and a second computing element to execute in lockstep in a fault-tolerant server, the apparatus comprising:
- (a) means for establishing communication between the first computing element and a communications link;
 - (b) means for establishing communication between the second computing element and the communications link;
 - (c) means for transmitting, by the first computing element, a first instruction to the communications link;
 - (d) means for transmitting, by the second computing element, a second instruction to the communications link; and
 - (e) means for comparing, by at least one of a local input-output (I/O) subsystem of the first computing element and a local I/O subsystem of the second computing element, the first instruction and the second instruction and indicating a fault of at least one of the first computing element and the second computing element in response thereto.